

Maximin Safety: When Failing to Lose is Preferable to Trying to Win*

Brad Gulko Samantha Leung
 Computer Science Department
 Cornell University
 Ithaca, NY 14853
 bgulko | samlyy@cs.cornell.edu

January 22, 2015

Abstract

We present a new decision rule, *maximin safety*, that seeks to maintain a large margin from the worst outcome, in much the same way minimax regret seeks to minimize distance from the best. We argue that maximin safety is valuable both descriptively and normatively. Descriptively, maximin safety explains the well-known *decoy effect*, in which the introduction of a dominated option changes preferences among the other options. Normatively, we provide an axiomatization that characterizes preferences induced by maximin safety, and show that maximin safety shares much of the same behavioral basis with minimax regret.

1 Introduction

Representing uncertainty using a probability distribution, and making decisions by maximizing expected utility, is widely accepted, founded on formal mathematical principles, and satisfies intuitive notions of rationality such as independence of irrelevant alternatives and the sure thing principle [20]. However, enforcing seemingly appealing concepts of rationality can ultimately lead to decisions inconsistent with what real humans consider reasonable. For example, observed behavior under unquantified (Knightian [14]/ strict [15]) uncertainty, such as that in the Ellsberg paradox [8], demonstrates how appealing concepts of rationality can lead to inconsistency with human choices. Alternative decision rules, such as maximin utility [25] and minimax regret [20, 17] provide rationally plausible decisions in ambiguous situations and can be used to resolve

*We thank Joseph Y. Halpern for useful discussions and anonymous reviewers for useful comments. Work partly supported by NSF grants IIS-0812045 and CCF-1214844, and ARO grant W911NF-09-1-0281.

such paradoxes, but still fail to explain some human behavioral patterns. A particularly illustrative example of such behavior is called the *decoy effect* [13], in which the introduction of a *dominated* option changes the preference among the *undominated* ones. While the decoy effect has been investigated in the psychology [6, 26] and economics literature [3, 22], we are unaware of any axiomatic treatment of it. To address this, we introduce a criterion called *safety* as the basis for a *maximin safety* decision rule.¹ Safety serves as a dual to regret that quantifies distance from a worst outcome, much as regret quantifies proximity to a best outcome. Maximin safety also satisfies familiar properties common to maximin utility and minimax regret, and hence also resolves the Ellsberg paradox. Moreover, maximin safety accommodates observed preferences that are incompatible with minimax regret and maximin utility. We demonstrate how safety-seeking behavior can produce the decoy effect, and show how maximin safety can explain it. We also extend Stoye’s [24] axiomatizations of standard decision rules to include maximin safety, thus allowing a comparison between maximin safety and state-of-the-art decision rules.

1.1 Relative Preferences and Regret

It is not hard to imagine situations in which performance *relative* to other possible outcomes is more important than *absolute* performance. Consider, for example, a group of duck hunters surprised by a hungry bear [5, 4]. The hunters all attempt to escape by running in the same direction while the slowest one despairs: “this is hopeless, we can never outrun the bear.” The hunter in front of him snickers, “I don’t need to outrun the bear, I just need to outrun *you*.” Whether the prospect is being picked from a group of peers for a date [2], winning a gold medal, or obtaining an ‘A’ in a class, success is often measured by relative performance, rather than by an absolute standard. One such preference for relative performance is embodied in the well-known decision theoretical concept of *regret* [20, 17]. While psychological literature on regret focuses on the bad feelings that occur *after* a choice leads to an inferior outcome, some also considers that anticipation of such negative emotions may influence the choice itself [22, 16, 19].

In this paper, we assume that uncertainty is captured by a set of possible worlds, one of which is the true state of the world. Regret is a measure of distance between the value of a considered outcome and the value of the best possible outcome, under a given state. This leads to an important property that is always true for regret – the introduction of a *dominated* option does not change the regrets of the existing options. We will refer to this property as *independence of dominated alternatives* (IDA). Those who believe in regret avoidance may think that this property is perfectly reasonable. For example, suppose you have a \$10 bill and you can either buy a \$10 lottery ticket, or two \$5 lottery tickets. Most would agree that your choice should not be affected

¹This decision rule has been mentioned in passing, inside a proof by Hayashi [11], where it was referred to as ‘maximin joy’. We use the term ‘safety’ rather than ‘joy’ to avoid confusion with the concept called ‘joy of winning’ in [11].

by a dominated third option, “burning the \$10 bill”. Other standard decision rules, such as expected utility maximization, have even stronger independence guarantees. The ranking of two choices under expected utility maximization is *menu-independent*, i.e., completely independent of the set of feasible choices (the *menu*). Menu-independence implies IDA. In contrast, regret-based preferences are menu-dependent, but since they conform to IDA, they are not compatible with observed biases sensitive to dominated options [2]. While IDA seems intuitively appealing, there is a great deal of empirical evidence that human preferences are indeed affected by dominated options in measurable and sometimes profound ways.

1.2 The Decoy Effect in Decision Theory

Suppose you are offered \$6 in cash, and the option of trading it for a Cross pen. The pen is nice, but you have plenty of pens, so decide to keep the cash. Right before you walk away, you are offered an alternative pen in exchange for the \$6. You see the new pen and find it hideous. A smile comes to your face as you turn around and say, “you know, I’ll take that original Cross pen after all.”

This story dramatizes an actual experiment [13]. When the first choice was offered to 106 people, 64% took the cash, 36% took the pen. When the second pen was added to the offer to 115 other subjects, 52% took the cash, 46% took the Cross pen, and 2% took the decoy. Generally, a *decoy* is an option that is designed to be inferior to another option in every way (i.e., it is a dominated option). Despite the intuitive appeal of IDA, the presence of a dominated option drove selection of the Cross pen from 36% to 46%. In this paper, we focus on a particular class of *decoy effect*, called *asymmetric dominance*, which occurs when the decoy is dominated by one existing alternative, but not by another. Empirical studies show that the decoy is rarely chosen, but its addition to a set of choices consistently drives decision makers toward the *dominating choice*.

Numerous empirical studies have also shown decoy effects in class action settlements [27], recreational land management [3], choice of healthcare plans and political candidates [12], purchase of consumer goods such as cameras and personal computers [22], restaurant choices [13], and even romantic attraction [2]. Surprisingly, a decoy effect can occur even if the decoy is not actually an option, but merely a recent memory of an option (a phantom decoy [9, 6]). Furthermore, the decoy effect is not limited to humans, but is also observed in honeybees and grey jays [21].

In an attempt to explain the decoy effect, experts in the behavioral sciences have offered a variety of domain-specific analyses, including “perceptual framing” [13], “value-shift” [26], “extremeness aversion” [22], and “contrast bias” [22, 27]. All of these explanations focus on valuing the discrepancy between the decoy and the dominating alternative. Intuitively, this provides a compelling example of preferring the margin of safety from the worst outcome. As we are not aware of any formalization in decision theory that is consistent with the intuitive preference for “margin of safety”, we offer one here.

To illustrate our new decision rule, recall the example of the unfortunate duck hunters. As they run from the bear, they approach a blind curve and have no idea what is around it: it could be wet or dry. If it is dry they will cover the most ground if they try to run faster, however if it is wet (thus slippery) they will be better off if they slow down and maintain balance. The options and the distance traveled under each circumstance are summarized in Table 1. In general, exerting excessive effort on a wet road leads to slipping and less distance covered; exerting effort on a dry road leads to more distance covered.

	Wet Road	Dry Road
Sprint	1	9
Hustle	3	6
Jog	2	2

Table 1: Hunters running from a bear.

If the probability of the road conditions is unknown, and only the first two options are available (sprint and hustle), there is no intuitively preferred choice and we may assume there are enough hunters such that at least one will pick each option. However, if we add a new option, jog, something interesting happens. As jog is dominated by hustle, IDA requires that its availability should not change the preferences among the other options. However, regardless of whether the road is wet or dry, hustle is never the worst alternative: if the road is wet, hustle (3) is faster than sprint (1), and if the road is dry, hustle (6) is faster than jog (2). In either case, selecting hustle prevents the hunter from being the slowest and getting caught by the bear.

While it may be callous, it seems perfectly reasonable for a hunter to decide to run just fast enough to make sure there is someone behind him. In other words, the most sensible decision might be to run just fast enough to guarantee the maximum possible margin between himself and the slowest runner, in the worst scenario. This margin between the hunter and his slowest compatriot can be considered a measure of *safety*, which is at the heart of our paper.

The rest of the paper proceeds as follows. Section 2 provides a formalization of the decoy paradox along with basic decision-theoretical notation. Section 3 describes the relationship between minimax regret and *maximin safety* and shows how maximin safety resolves the decoy paradox. Section 4 provides an axiomatic characterization of maximin safety. Section 5 suggests a unification of utility, regret, and safety using *anchoring functions*, and also considers a generalization to qualitative relative preferences.

2 The Formal Framework

Given a set S of *states* and a set X of *outcomes*, an *act* a (over S and X) is a function mapping S to X . The set of all acts is thus X^S , which we will denote by A . For simplicity in this paper, we take S to be finite. Associated with each outcome $x \in X$ is a *utility*: $U(x)$ is the utility of outcome x . For convenience, we will omit the explicit representation of the outcome, and denote $U(a(s))$ by $U(a, s)$ for each state $s \in S$. We call a tuple (S, X, U) a (non-probabilistic) decision problem. To define regret and safety, we need to assume that we are also given a set $M \subseteq A$ of feasible acts, called the *menu*. The reason for the

menu is that, as we have shown, regret and safety can depend on the menu. We will only consider finite menus, from which randomized strategies can be chosen.

Consider the problem of a decision maker (DM) contemplating a camera purchase, summarized in Table 2. The DM has a choice between buying a rugged travel camera (a_1) that takes decent pictures in a wide variety of circumstances, and buying a delicate sports camera with higher speed and image quality (a_2). Each state characterizes the possible situations that a purchaser may experience during the useful life of the camera (Will the DM experience harsh conditions? Or win tickets to the World Cup?) The utility $U(a, s)$ of act a under state s represents an abstract net value to the DM if the true world is state s .

	s_1 Safari	s_2 World Cup
a_1 : Travel	4	4
a_2 : Sports	2	6
a_3 : Decoy	3	3

Table 2: Utilities in the camera purchase example.

If the DM ends up going on a safari (s_1), then act a_1 results in moderate quality pictures of exciting wildlife ($U(a_1, s_1) = 4$), but act a_2 results in a few exquisite shots and many missed opportunities ($U(a_2, s_1) = 2$). On the other hand, if the DM goes to the World Cup (s_2), then act a_2 results in many great pictures in a safe environment ($U(a_2, s_2) = 6$), while act a_1 provides only moderate quality pictures ($U(a_1, s_2) = 4$).

If the DM can assign probabilities $P(s_1)$ and $P(s_2)$ to the states, she can calculate an expected utility $E[U(a_i)] = \sum_{s \in S} P(s)U(a_i, s)$, and simply select the act that maximizes expected utility. However, if the state probabilities are unavailable, we have unquantified uncertainty. In such cases, the DM must find another method for aggregating the utility of each act across states in order to assign a *value* to each camera. Here we will focus on the methods of maximax utility, minimax utility, and minimax regret. To understand minimax regret, we need to define the notion of regret. For a menu M and act $a \in M$, the regret of a with respect to M and decision problem (S, X, U) is

$$\max_{s \in S} (\max_{a' \in M} U(a', s) - U(a, s)).$$

We denote this as $reg_M^{(S, X, U)}(a)$, and usually omit the superscript (S, X, U) .

When comparing decision rules, it is often convenient to define a *value function* that assigns a numeric value to each act, for the purpose of ranking the acts. Formally, for a decision problem (S, X, U) , a value function is a function

$$V^{(S, X, U)}(a, M) : X^S \times 2^A \rightarrow \mathbb{R}.$$

We will usually omit the superscript (S, X, U) and just write $V(a, M)$, or $V(a)$ if the value function is menu-independent.

Decision Rule	Value of an act a	Decision rule description	Best
maximax utility	$V(a) = \max_{s \in S} U(a, s)$	Optimize the best-case outcome.	a_2
maximin utility	$V(a) = \min_{s \in S} U(a, s)$	Optimize the worst-case outcome.	a_1
minimax regret	$V(a, M) = -\text{reg}_M(a)$	Pick an act to minimize the worst-case distance from the best outcome.	a_1, a_2

Table 3: Standard decision rules and most valued acts in the camera example.

We say that the value function V *represents* the family of preference relations $\succ_{V,M}$, if for all menus M and all $a, a' \in M$,

$$a \succ_{V,M} a' \Leftrightarrow V(a, M) > V(a', M).$$

In other words, act a is (strictly) preferred to act a' with respect to menu M if and only if $V(a, M) > V(a', M)$. The value functions and preferences of several standard decision rules are given in Table 3.

Now, perhaps the camera vendor would like to sell more travel cameras, so the vender puts an obsolete travel camera a_3 next to a_1 as a *decoy*. Camera a_3 has the same price as a_1 , but fewer features and lower picture quality. The vendor hopes to make a_1 more appealing by contrast with a_3 . Table 2 illustrates the decision problem when a_3 is added to the menu. The ranking between a_1 and a_2 according to each of the decision rules in Table 3 is unaffected by the introduction of a_3 to the menu. The addition of a_3 also illustrates the concept of dominance. We say that an act a *dominates* a' , if for all $s \in S$, $U(a, s) > U(a', s)$.

3 Maximin Safety

While minimax regret seeks to minimize separation from best outcomes, *maximin safety* is a conceptual dual that seeks to maximize separation from the worst outcomes. For a menu M and act $a \in M$, the safety of a in state s is defined as:

$$\text{safety}_M^{(S,X,U)}(a, s) = U(a, s) - \min_{a' \in M} (U(a', s)),$$

and in keeping with the convention for regret, the safety of an act is defined as:

$$\text{safety}_M^{(S,X,U)}(a) = \min_{s \in S} (\text{safety}_M^{(S,X,U)}(a, s)).$$

We will often omit the superscript (S, X, U) .

The family of maximin safety preferences $\succ_{\text{saf},M}$ represented by the *safety* value function satisfies, for all M and $a, a' \in M$,

$$a \succ_{\text{saf},M} a' \Leftrightarrow \text{safety}_M(a) > \text{safety}_M(a').$$

Now we reconsider the camera example using safety (Table 4). Without the decoy, both acts have the same safety of 0, since each act has the lowest utility in some state; so there is no clear safety preference. However, when the

	Utility		Safety (no decoy)		Safety (w. decoy)	
	s_1	s_2	s_1	s_2	s_1	s_2
a_1 : travel	4	4	2	0	2	1
a_2 : sports	2	6	0	2	0	3
a_3 : decoy	3	3			1	0

Table 4: Camera purchase with and without decoy.

	Utility		Regret		Safety		Optimal for
	s_1	s_2	s_1	s_2	s_1	s_2	
a_1	1	9	3	0	0	5	maximax utility
a_2	3	6	1	3	2	2	maximin safety
a_3	2	7	2	2	1	3	minimax regret
a_4	4	4	0	5	3	0	maximin utility

Table 5: Different decision rules select different acts for the same problem.

decoy is present, the act a_1 never has the lowest utility at any state, and thus it has a strictly positive safety. In this case, $safety_{\{a_1, a_2, a_3\}}(a_1)$ is the unique maximum among the acts $\{a_1, a_2, a_3\}$, and therefore a_1 is the preferred choice. The relative increase in the safety of an act due to the addition of the dominated act is an essential element in solving the decoy paradox. Intuitively, this may correspond to a sense that even if a particular act gets low utility in the realized state, the DM may think that “I’m better off than the fools who bought the worse camera”, or in a more positive light, “I must be getting a steal with this better camera for the same price”. In competitive survival games (such as the reality game show Survivor), the notion of maximizing safety may also embody a preference to maintain a maximal distance from the lowest performer, which reduces the chance of elimination. Table 5 compactly demonstrates how choices based on maximin safety differs from the other standard decision rules.

In the camera example, the addition of a decoy created a strict preference between two acts that were initially tied. The introduction of a dominated act can actually *reverse* preferences between acts. Table 6 shows a menu of three acts: $M = \{a_1, a_2, a_3\}$. Act a_2 has a minimum

	Utility			Safety(no decoy)			Safety (w. decoy a_4)		
	s_1	s_2	s_3	s_1	s_2	s_3	s_1	s_2	s_3
a_1	9	2	6	5	0	0	8	0	0
a_2	5	3	7	1	1	1	4	1	1
a_3	4	8	8	0	6	2	3	6	2
a_4	1	5	6				0	3	0

Table 6: Without a_4 , $M = \{a_1, a_2, a_3\}$, and $a_2 \succ_{saf, M} a_3$. Adding a_4 (dominated by a_3) reverses the maximin safety preference between a_2 and a_3 .

safety of 1, while both a_1 and a_3 have the lowest utility for some state, so each has minimum safety of 0. Consequently, a_2 is the most preferred choice under the safety preference. When a new choice a_4 is added, act a_4 is dominated by a_3 , but it has higher utility than the other acts in some states. This situation is known as *asymmetric dominance*, which is typically associated with decoy effects. In this example, asymmetric dominance guarantees that a_3 is never one of the worst choices, and thus has a strictly positive safety value. In other words, the addition of a_4 to the menu M does not affect the safety of a_1 or a_2 , but increases the safety of a_3 to make $a_3 \succ_{saf, M \cup \{a_4\}} a_2$.

4 Axiomatic Analysis

To provide an axiomatic characterization of maximin safety, we employ the standard *Anscombe-Aumann* (AA) framework [1], where outcomes are restricted to lotteries. Maximin safety is characterized by modifying one of the axioms in an existing characterization of minimax regret provided by Stoye [24].

Given a set Y of *prizes*, a *lottery* over Y is just a probability with finite support on Y . As in the AA framework, we let the set of *outcomes* be $\Delta(Y)$, the set of all lotteries over Y . Thus, *acts* are functions from S to $\Delta(Y)$. We can think of a lottery as modeling objective, quantified uncertainty, while the states model unquantified uncertainty. The technical advantage of considering such a set of outcomes is that we can consider convex combinations of acts. If f and g are acts, define the act $\alpha f + (1 - \alpha)g$ to be the act that maps a state s to the lottery $\alpha f(s) + (1 - \alpha)g(s)$. For simplicity, we follow Stoye [24] and restrict to menus that are the convex hull of a finite number of acts, so that if f and g are acts in M , then so is $pf + (1 - p)g$ for all $p \in [0, 1]$.

In this setting, we assume that there is a utility function U on prizes in Y , and that there are at least two prizes y_1 and y_2 in Y , with different utilities. Note that $l(y)$ is the probability of getting prize y if lottery l is played. We will use l^* to denote a constant act that maps all states to l . The utility of a lottery l is just the expected utility of the prizes obtained, that is, $u(l) = \sum_{\{y \in Y: l(y) > 0\}} l(y)U(y)$. The expected utility of an act f with respect to a probability \Pr is then just $u(f) = \sum_{s \in S} \Pr(s)u(f(s))$, as usual. Given a set Y of prizes, a utility U on the prizes, and a state space S , we have a family $\succeq_{saf, M}^{S, \Delta(Y), u}$ of preference orders on acts determined by maximin safety, where u is the utility function on lotteries as determined by U .² For convenience, from here on we will write $\succeq_M^{S, Y, U}$ rather than $\succeq_{saf, M}^{S, \Delta(Y), u}$. We will state the axioms in a way such that they can be compared to standard axioms and those for minimax regret in [24]. The axioms are universally quantified over acts f, g , and h , menus M and M' , and $p \in (0, 1)$. Whenever we write $f \succeq_M g$ we assume that $f, g \in M$.

² We let $f \succeq_{saf, M}^{S, \Delta(Y), u} g$ iff $g \not\succeq_{saf, M}^{S, \Delta(Y), u} f$, and $f \sim_M g$ iff $f \succeq_M g$ and $g \succeq_M f$.

Axiom 1. (*Monotonicity*) $f \succeq_M g$ if $(f(s))^* \succeq_{\{(f(s))^*, (g(s))^*\}} (g(s))^*, \forall s \in S$.

Axiom 2. (*Completeness*) $f \succeq_M g$ or $g \succeq_M f$.

Axiom 3. (*Nontriviality*) $f \succ_M g$ for some acts f and g and menu M .

Axiom 4. (*Mixture Continuity*) If $f \succ_M g \succ_M h$, then there exists $q, r \in (0, 1)$ such that $qf + (1 - q)h \succ_M g \succ_M rf + (1 - r)h$.

Axiom 5. (*Transitivity*) $f \succeq_M g \succeq_M h \Rightarrow f \succeq_M h$.

Menu-independent versions of Axioms 1 to 5 are standard in other axiomatizations, and in particular hold for maximin utility. Axiom 3 is used in the standard axiomatizations to get a nonconstant utility function in the representation. While maximin safety does not satisfy menu-independence, it does satisfy menu-independence when restricted to menus consisting of only constant acts. This property is captured by the following axiom.

Axiom 6. (*Menu independence for constant acts*) If l^* and $(l')^*$ are constant acts, then $l^* \succeq_M (l')^*$ iff $l^* \succeq_{M'} (l')^*$.

We also have a menu-dependent version of the von Neumann-Morgenstern (VNM) Independence axiom. Like the VNM Independence axiom, Axiom 7 says that ranking between two acts does not change when both acts are mixed with a third act; but unlike VNM Independence, the menu used to compare the original acts in Axiom 7 is different from that used to compare the mixtures. Axiom 7 holds for minimax regret and maximin safety, but not for maximin utility.

Axiom 7. (*Independence*) $f \succeq_M g \Leftrightarrow pf + (1 - p)h \succeq_{pM + (1 - p)h} pg + (1 - p)h$.

Axiom 8. (*Symmetry*) For a menu M , suppose that $E, F \in 2^S \setminus \{\emptyset\}$ are disjoint events such that for all $f \in M$, f is constant on E and on F . Define f' by

$$f'(s) = \begin{cases} f(s') \text{ for some } s' \in E, & \text{if } s \in F \\ f(s') \text{ for some } s' \in F, & \text{if } s \in E \\ f(s) & \text{otherwise} \end{cases}$$

Let M' be the menu generated by replacing every act $f \in M$ with f' . Then

$$f \succeq_M g \Leftrightarrow f' \succeq_{M'} g'$$

Symmetry, which is one of the characterizing axioms for minimax regret in [24], captures the intuition that no state can be considered more or less likely than another. Therefore Symmetry helps distinguish the probability-free decision rules maximin utility, minimax regret, and maximin safety, from their probabilistic counterparts [10, 23].

Axiom 9. (*Ambiguity Aversion*) $f \sim_M g \Rightarrow pf + (1 - p)g \succeq_M g$.

Axiom 9 says that the decision maker weakly prefers to hedge her bets. Axioms 1-9 are all part of the characterization in [24] of minimax regret (which consists of Axioms 1-9 and Symmetry). Axioms 1-5 and 9 are also sound for the maximin decision rule [24].

In [24], one of the axioms characterizing minimax regret is Independence of Never Strictly Optimal alternatives (INA), which states that adding or removing acts that are not strictly potentially optimal in the menu does not affect the ordering of acts.³ By varying this INA axiom, we obtain a characterization for maximin safety. We say that an act a is *never strictly worst relative to M* if, for all states $s \in S$, there is some $a' \in M$ such that $a(s) \succeq a'(s)$.

Axiom 10. (*Independence of Never Strictly Worst Alternatives (INWA)*) If an act a is never strictly worst relative to M , then $f \succeq_M g$ iff $f \succeq_{M \cup \{a\}} g$.

Although adding acts to the menu, in general, can affect minimax regret preferences, INA implies the Independence of Dominated Alternatives property that we used earlier when discussing the decoy effect. Thus, INA guarantees that minimax regret can never be compatible with the decoy effect.

Theorem 1. For all Y, U, S , the family of maximin safety preference orders $\succeq_{saf, M}^{S, Y, U}$ induced by a decision problem $(S, \Delta(Y), u)$ satisfies Axioms 1–10. Conversely, if the family of preference orders \succeq_M on the acts in $\Delta(Y)^S$ satisfies Axioms 1–10, then there exists a utility function U on Y that determines a utility u on $\Delta(Y)$ such that $\succeq_M = \succeq_{saf, M}^{S, Y, u}$. Moreover, U is unique up to affine transformations.

Proof. The soundness of the axioms are straightforwardly verified, so we show only the completeness of the axioms. We will use the same general sequence of arguments that Stoye uses in [24]. First, we establish a nonconstant utility function U , where constant acts are ranked by their expected utilities. Since we have the standard axioms (1 – 5), we get U from standard arguments, and it is unique up to affine transformations. Next, we observe the following lemma:

Lemma 1. Suppose the family \succeq_M satisfies Axioms 1-10, and \succeq_{M^+} is representable by maximin safety, where M^+ is the menu of all acts with nonnegative utilities. Then the family \succeq_M is representable by maximin safety.

Lemma 1 follows from an argument analogous to that for regret in [24]. The next step is to establish that the axioms on \succeq_M restrict \succeq_{M^+} to satisfy the axioms of ambiguity aversion, monotonicity, completeness, transitivity, non-triviality, and symmetry. It is a straightforward verification that will not be reproduced here. Theorem 1 (iii) of [24] then implies that \succeq_{M^+} is the maximin utility ordering. Next, let g_M be an act such that $u \circ g_M(s) = -\min_{h \in M} u(h, s)$, so that we have

³ An act h is *never strictly optimal relative to M* if, for all states $s \in S$, there is some $f \in M$ such that $(f(s))^* \succeq (h(s))^*$.

$$\begin{aligned}
f \succeq_M g &\Leftrightarrow \frac{1}{2}f + \frac{1}{2}g_M \succeq_{M+} \frac{1}{2}g + \frac{1}{2}g_M \\
&\Leftrightarrow \min_{s \in S} u(\frac{1}{2}f + \frac{1}{2}g_M, s) \geq \min_{s \in S} u(\frac{1}{2}g + \frac{1}{2}g_M, s) \\
&\Leftrightarrow \min_{s \in S} (\frac{1}{2}(u(f, s) - \min_{h \in M} u(h, s))) \geq \min_{s \in S} (\frac{1}{2}(u(g, s) - \min_{h \in M} u(h, s))).
\end{aligned}$$

□

□

The characterizing axioms serve as a justification for maximin safety in the sense that behaving as a safety maximizer is equivalent to accepting the axioms. Axioms 1-7 are standard and broadly accepted to be reasonable, while symmetry and ambiguity aversion are implied by both maximin utility and minimax regret. Whether the INA axiom (for regret) or the INWA axiom (for safety) is more reasonable would depend on the individual and the nature of the decision problem. Thus, we believe that the reasonableness of the maximin safety decision rule is comparable to that of minimax regret.

Individual necessity of the axioms can be established, as is commonly done [11, 24], by giving examples of preferences that satisfy all the axioms except for the one whose necessity is being shown. For the axioms shared with minimax regret, the same examples found in [24] shows their individual necessity. For the INWA axiom, the required example is minimax regret. Indeed, a decision rule equivalent to maximin safety was used by Hayashi [11] as an example to justify minimax regret's entailment of INA.

Clearly, just as minimax regret is readily generalized to *minimax expected regret* when uncertainty is represented by a set of probability distributions over the state space, maximin safety can be readily extended to *maximin expected safety* in the same manner. As one would expect, given an axiomatization of minimax expected regret [23], the modification of the INA axiom to INWA results in an axiomatization for maximin *expected* safety.

5 Discussion, Generalizations, and Future Work

Both minimax regret and maximin safety embody preferences based on *relative*, rather than *absolute* utility. In Table 5, the act preferred by safety has a lower minimum utility than the act preferred by maximin utility, just as the act picked by minimax regret neglects a higher maximum utility in order to minimize the *margin* to each state's maximum utility. The shared preference for relative over absolute performance is reflected in a striking similarity in the structure of the value functions for regret and safety. In comparison, minimax regret can be expressed for all acts a, b as:

$$a \succ_{reg, M} b \text{ iff } \min_{s \in S} (U(a, s) - \max_{a' \in M} U(a', s)) > \min_{s \in S} (U(b, s) - \max_{a' \in M} U(a', s)).$$

Similarly, maximin safety is represented for all acts a, b as

$$a \succ_{saf, M} b \text{ iff } \min_{s \in S} (U(a, s) - \min_{a' \in M} U(a', s)) > \min_{s \in S} (U(b, s) - \min_{a' \in M} U(a', s)).$$

The structural resemblance suggests a common form for the value function. By defining a menu-dependent *anchoring function* $t : S \times 2^A \rightarrow \mathbb{R}$, we can represent several previously discussed value functions as:

$$V_t(a, M) = \min_{s \in S} U'(a, s, M, t),$$

where $U'(a, s, M, t) = U(a, s) - t(s, M)$ can be viewed as an *anchored effective utility*. One can see that V_t represents maximin utility if $t(s, M) = 0$; minimax regret if $t(s, M) = \max_{a' \in M} U(a', s)$; and maximin safety if $t(s, M) = \min_{a' \in M} U(a', s)$. Note that by varying just the anchoring function t , we can obtain all the mentioned decision rules, and more. While we focus only on maximin safety in this paper, other forms for $t(s, M)$ maximize the positive margin from a state-dependent average, median, or some other characteristic of interest to a DM. For example, college students might seek to conservatively maximize their margin above a desired quantile, in order to achieve a particular grade.

The present work is motivated by behavioral observations of the decoy effect that are typically described in empirical quantities such as distance, price and volume, and thus is most intuitive in a *quantitative* framework. However, the key observation is that safety, like regret, is a notion of relative performance with respect to a set of outcomes, rather than absolute performance. As absolute quantitative utility $U : X \rightarrow \mathbb{R}$ can be generalized to a qualitative framework by replacing the U with a mapping $X \rightarrow L$ for some ordered set L , relative utility may be made qualitative by considering the mapping with $2^X \times X \rightarrow L$. In the case of safety and regret, the particular element of 2^X is the set of all possible outcomes in a state, given a menu of acts. Aggregation of N state-specific orderings into an ordering over acts can be accomplished by an aggregation function $M : L^N \rightarrow L$ [18]. This generalization can be readily applied to various characterizations of uncertainty, including probability, plausibility, and the strict uncertainty used in this paper [15]. While the authors expect that the present quantitative axiomatization can be adapted to a qualitative framework (see, e.g. [7]), it is beyond the scope of the current paper.

References

- [1] F. Anscombe and R. Aumann. A definition of subjective probability. *Annals of Mathematical Statistics*, 34:199–205, 1963.
- [2] D. Ariely. *Predictably Irrational: The Hidden Forces That Shape Our Decisions*. HarperCollins, 1 edition, Feb. 2008.
- [3] I. J. Bateman, A. Munro, and G. L. Poe. Decoy effects in choice experiments and contingent valuation: Asymmetric dominance. *Land Economics*, 84(1), 2008.
- [4] L. Chao, K. A. Hanley, C. L. Burch, C. Dahlberg, and P. E. Turner. Kin selection and parasite evolution: Higher and lower virulence with hard and soft selection. *The Quarterly Review of Biology*, 75(3):pp. 261–275, 2000.

- [5] M. Crawley. Some sociological misconceptions, center for population biology. *Pan-European Conference on the Potential Long-Term Ecological Impact of Genetically Modified Organisms: Proceedings*, November 1993:276, 1993.
- [6] J. R. Doyle, D. J. O'Connor, G. M. Reynolds, and P. A. Bottomley. The Robustness of the Asymmetrically Dominated Effect: Buying Frames, Phantom Alternatives, and In-Store Purchases. *Psychology and Marketing*, 16(3), 1999.
- [7] D. Dubois, H. Fargier, and P. Perny. Qualitative decision theory with preference relations and comparative uncertainty: An axiomatic approach. *Artificial Intelligence*, 148(12):219 – 260, 2003.
- [8] D. Ellsberg. Risk, ambiguity, and the savage axioms. *QJE*, 75(4), 1961.
- [9] P. H. Farquhar and A. R. Pratkanis. Decision structuring with phantom alternatives. *Management Science*, 39(10):pp. 1214–1226, 1993.
- [10] I. Gilboa and D. Schmeidler. Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics*, 18(2):141–153, April 1989.
- [11] T. Hayashi. Regret aversion and opportunity dependence. *JET*, 139(1), 2008.
- [12] W. Hedgcock, A. R. Rao, and H. A. Chen. Could Ralph Nader’s Entrance and Exit Have Helped Al Gore? The Impact of Decoy Dynamics on Consumer Choice. *Journal of Marketing Research*, 46(3), 1999.
- [13] J. Huber, J. W. Payne, and C. Puto. Adding asymmetrically dominated alternatives: Violations of regularity and the similarity hypothesis. *JCR*, 9(1), 1982.
- [14] F. H. Knight. *Risk, Uncertainty and Profit*. Houghton Mifflin, 1921.
- [15] R. B. Larbi, S. Konieczny, and P. Marquis. A characterization of optimality criteria for decision making under complete ignorance. In *KR*, 2010.
- [16] R. P. Larrick and T. L. Boles. Avoiding regret in decisions with feedback: A negotiation example. *OBHDP*, 63(1):87–97, July 1995.
- [17] R. D. Luce and H. Raiffa. *Games and decisions : introduction and critical survey*. Wiley, 1957.
- [18] J.-L. Marichal. An axiomatic approach of the discrete Sugeno integral as a tool to aggregate interacting criteria in a qualitative framework. *IEEE Transactions on Fuzzy Systems*, 9(1):164–172, 2001.
- [19] I. Ritov. Probability of regret: Anticipation of uncertainty resolution in choice. *Organizational Behavior and Human Decision Processes*, 66(2):228–236, May 1996.

- [20] L. J. Savage. *The Foundations of Statistics*. John Wiley, 1954.
- [21] S. Shafir, T. A. Waite, and B. H. Smith. Context-Dependent violations of rational choice in honeybees (*apis mellifera*) and gray jays (*perisoreus canadensis*). *Behavioral Ecology and Sociobiology*, 51(2):180–187, 2002.
- [22] I. Simonson and A. Tversky. Choice in context: Tradeoff contrast and extremeness aversion. *Journal of Marketing Research*, 29(3):281–295, 1992.
- [23] J. Stoye. Axioms for minimax regret choice correspondences. *JET*, 146(6), 2011.
- [24] J. Stoye. Statistical decisions under ambiguity. *Theory and Decision*, 70(2):129–148, February 2011.
- [25] A. Wald. *Statistical Decision Functions*. John Wiley, NY, 1950.
- [26] D. H. Wedell. Distinguishing among models of contextually induced preference reversals. *Journal of Experimental Psychology*, 17(4):767–778+, 1991.
- [27] A. S. Zimmerman. Funding Irrationality. *Duke Law Journal*, 59(1105), 2010.